CHARLES PARKER

658 Rockland Avenue, Lake Bluff, IL 60044 • 253.229.4622 • clp@clparker.org

EDUCATION

Ph.D., Computer Science September 2002 - August 2007
Oregon State University *Corvallis, OR*
Graduate Research: Machine Learning

B.S., Electrical and Computer Engineering September 1998 - December 2001
Northwestern University *Evanston, IL*
Curriculum Focus: Microcontroller Systems Design

WORK EXPERIENCE

Vice President, Machine Learning Algorithms November 2015 - Present
BigML, Inc. *Corvallis, OR*
Planned and developed extensions to a "software as a service" platform for machine learning. Collaborated with multi-disciplinary team to define goals and directions for future development. Supported customers and potential investors with product exploration and early utilization. Gave general presentations on machine learning to assist with marketing and education initiatives.

Worked as developer on the following projects:

Bayesian Parameter Optimization Implemented several popular algorithms for automatic parameter optimization using an in-house domain-specific language, which enabled users optimize the hyperparameters of machine learning algorithms with little expertise and minimal effort.

Basic and Advanced Natural Language Processing Served as primary developer for all natural language processing features on the platform, including basic textual processing such as tokenization and word stemming through higher-level techniques such as latent topic modeling.

Deep Learning Server Architecture Engineered a lightweight, python-based server for learning deep neural networks. Utilized both TensorFlow and Theano to enable high-speed multicore computation and GPU utilization.

Senior Quantitative Research Analyst November 2012 - November 2015
Allston Trading *Chicago, IL*
Conducted research and analysis to support and extend a portfolio of futures trading strategies. Utilized machine learning techniques to build a generalized strategy creation framework for market making strategies across a variety of U.S. and European futures, including equity index, treasury, currency and energy products. Parallelized analysis using a MapReduce framework on Hadoop. Developed web-based tools and visualizations for post-trade analysis using JavaScript and D3.

Senior Machine Learning Engineer / Strategic Adviser September 2011 - November 2015
BigML, Inc. *Corvallis, OR*
Collaborated with development team to build a scalable infrastructure for machine learning on large data. Utilized tools such Hadoop and Storm to specify a generic interface from machine learning algorithms to massively parallel architectures. Developed performance benchmarks, data generators, and algorithms for data set manipulation. Worked closely with clients, potential clients, and investors to help understand and utilize the system. Created short films and animations to be utilized as product demonstrations.

Adjunct Faculty January 2009 - May 2011
University of Rochester *Rochester, NY*
Taught courses in the computer science department at the University of Rochester. Participated in curriculum

committee meetings to design new courses, and implemented these designs. Collaborated with undergraduate researchers to design course syllabus and materials. Managed a staff of up to 15 undergraduate teaching assistants, graduate teaching assistants, and "workshop facilitators". Developed exams for student evaluation.

Taught the following courses:

- Introduction to Computer Programming
- Introduction to Computer Technology
- Multimedia Computer Applications
- Computing for the Humanities
- Computing for Engineering

Research Scientist September 2007 - September 2011
Eastman Kodak Company *Rochester, NY*
Engaged in industrial research for product development and improvement, especially in the fields of machine learning, data mining, computer vision, and natural language processing. Worked with business units and external partners to leverage and implement practical research findings for use in products. Developed patents to improve and expand Kodak's intellectual property portfolio.

Participated in the following research efforts:

Pattern mining for machine reliability Applied machine learning and pattern mining techniques to a database of observed part failures to determine superior methods for predicting failure prediction and correlations between part failure rates. Delivered technical report to business unit on findings, including recommendations for operator retraining based on data analysis.

Text mining and topic modeling for scanned document databases Used standard string processing techniques and latent Dirichlet allocation to determine logical groupings of a set of unstructured documents acquired by scanner. Used these results to prototype a system for unsupervised organization of document collections.

Audio event localization and categorization in consumer video Developed algorithms for locating audio events (e.g., applause, music, laughter) in consumer videos. Compared several different feature extraction methods and learning algorithms to classify audio frames. Utilized conditional random fields to account for sequential effects. Delivered research code suitable for audio indexing of videos or assisting in video summarization.

Automatic print dating and text source determination Served as machine learning expert for the Kodak "Scan The World" project. Developed a decision tree-based algorithm to identify the dates of photographic prints based on visually determined attributes of the given print. Developed another algorithm to distinguish machine from handwritten text on the back of photographic prints using edge direction and color information. Delivered product-ready code to consumer imaging business unit.

Graduate Research Assistant January 2006 - September 2007
Oregon State University *Corvallis, OR*
Developed and tested machine learning research code. Integrated with other machine learning modules from several different universities and the Advanced Technology Laboratory at Lockheed Martin. Participated in teleconferences and meetings to focus research direction and resolve integration issues. Prepared requirements documents and component walkthroughs to satisfy grant administration requirements.

Graduate Teaching Assistant September 2002 - June 2006
Oregon State University *Corvallis, OR*
Assisted students in comprehending course material. Assisted professors in developing class exercises and material. Created and delivered several lectures in professors' absence. Demonstrated problem solving techniques in group recitation sessions and through one-on-one tutoring.

Programmer/Analyst December 2001 - June 2002
Grossman and Associates *Chicago, IL*
Programmed, debugged, and tested modules within a complex financial accounting software program. Dealt with client issues and service requirements, constructing software fixes and new features as appropriate. Wrote end-user documentation for completed modules.

PUBLICATIONS

Copies of some of these publications are available at `http://www.clparker.org/`.

[1] Charles Parker. On Measuring the Performance of Binary Classifiers. *Knowledge and Information Systems*, 35:131–152, 2013.

[2] X. Shelley Zhang, S. Yoon, P. DiBona, D. S. Appling, L. Ding, J. R. Doppa, D. Greeny, J. K. Guo, U. Kuter, G. Levine, R. L. MacTavish, D. McFarlane, J. R. Michaelis, Hala Mostafa, S. Ontañón, C. Parker, J. Radhakrishnan, A. Rebguns, B. Shrestha, Z. Song, E. B. Trewhitt, Huzaifa Zafar, Chongjie Zhang, Daniel Corkill, G. DeJong, T. G. Dietterich, S. Kambhampati, and Victor Lesser. An Ensemble Architecture for Learning Complex Problem-Solving Techniques from Demonstration. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 4(3):75:1–75:38, 2012.

[3] Charles Parker. An analysis of performance measures for binary classification. In *The International Conference on Data Mining*, pages 517–526, Vancouver, Canada, December 2011.

[4] Charles Parker, Dhiraj Joshi, Phoury Lei, and Jiebo Luo. Finding geographically representative music via social media. In *Proceedings of the First International ACM Workshop on Music Information Retrieval with User-centered and Multimodal Strategies*, pages 27–32, November 2011.

[5] Charles Parker. Performance measure choices for evaluating binary classifiers. Technical Report 345598M, Eastman Kodak Company, Rochester, NY, January 2011.

[6] Charles Parker. An exploration of semantic audio classification. Technical Report 345596K, Eastman Kodak Company, Rochester, NY, December 2010.

[7] Charles Parker. Anchor point selection by KL-divergence. In *WNYIPW '10: The Western New York Image Processing Workshop*, Rochester, NY, November 2010.

[8] Charles Parker. An empirical study of feature extraction methods for audio classification. In *ICPR '10: The Twentieth International Conference on Pattern Recognition*, pages 4593–4596, Istanbul, Turkey, August 2010.

[9] Charles Parker and Paul Messier. Automating art print authentication using metric learning. In *IAAI '09: The Twenty-First Innovative Applications of Artificial Intelligence Conference*, Pasadena, CA, July 2009.

[10] X. Shelley Zhang, S. Yoon, P. DiBona, D. S. Appling, L. Ding, J. R. Doppa, D. Greeny, J. K. Guo, U. Kuter, G. Levine, R. L. MacTavish, D. McFarlane, J. R. Michaelis, Hala Mostafa, S. Ontañón, C. Parker, J. Radhakrishnan, A. Rebguns, B. Shrestha, Z. Song, E. B. Trewhitt, Huzaifa Zafar, Chongjie Zhang, Daniel Corkill, G. DeJong, T. G. Dietterich, S. Kambhampati, and Victor Lesser. An ensemble learning and problem solving architecture for airspace management. In *IAAI '09: The Twenty-First Innovative Applications of Artificial Intelligence Conference*, pages 203–210, Pasadena, CA, July 2009.

[11] Charles Parker. An analysis of Kodak nexpress failure data. Technical Report 344794L, Eastman Kodak Company, Rochester, NY, December 2008.

[12] Charles Parker. *Structured Gradient Boosting*. PhD thesis, Oregon State Unversity, Corvallis, OR, August 2007.

[13] Charles Parker, Prasad Tadepalli, Weng-Keen Wong, Thomas Dietterich, and Alan Fern. Learning from demonstrations via structured prediction. In *AAAI '07 Workshop on Acquiring Planning Knowledge via Demonstration*, Vancouver, BC, Canada, July 2007.

[14] Charles Parker, Alan Fern, and Prasad Tadepalli. Learning for efficient retrieval of structured data with noisy queries. In *ICML '07: The Twenty-Fourth International Conference on Machine Learning*, pages 729–736, Corvallis, OR, June 2007.

[15] Charles Parker, Alan Fern, and Prasad Tadepalli. Gradient boosting for sequence alignment. In *AAAI '06: The Twenty-First National Conference on Artificial Intelligence*, pages 452–457, Boston, MA, July 2006.

[16] Charles Parker. Applications of binary classification and adaptive boosting to the query-by-humming problem. In *International Symposium on Music Information Retrieval*, pages 245–251, London, England, September 2005.

[17] Charles Parker. Examining synthetic databases in melodic retrieval testing. In *International Conference on Computer Music*, Miami, FL, November 2004.

[18] Charles Parker. A fast tree-based method for melodic retrieval. In *ACM Joint International Conference on Digital Libraries*, Tucson, AZ, June 2004.

[19] Charles Parker. Towards intelligent string matching in query-by-humming systems. In *IEEE International Conference on Multimedia and Expo*, pages 25–28, Baltimore, MD, June 2003.

AWARDS AND HONORS

- Promoted to Eastman Kodak Research Associate, 2011
- Awarded the Oregon State University College of Engineering "Outstanding Research Assistant of the Year" award 2006-2007 academic year
- Elected Computer Science Graduate Student Association vice-president, 2005-2006 academic year
- Nominated for the Oregon State University College of Engineering "Outstanding Teaching Assistant of the Year" award 2002-2003 and 2004-2005 academic years
- Elected to the *Upsilon Pi Epsilon* honor society for the computing sciences, 2003-2004 academic year

INVITED TALKS

- MLPrague, "Really Automating Machine Learning" Prague, Czech Republic. March, 2018.
- PAPIs.io, "Automated Machine Learning: Mostly Unhelpful" Boston, MA. October 2017.
- International Joint Conference on Artificial Intelligence, "Automated Machine Learning: Mostly Unhelpful" Melbourne, Australia. August 2017.
- Real-Time Big Data Meetup, "Real-time Machine Learning". Menlo Park, CA. March, 2013.
- KDD BigMine-2012: First International Workshop on Big Data Mining, "Unexpected Challenges in Large Scale Machine Learning". August, 2012.
- Galois, Inc., "An Analysis of Analysis". April, 2012.
- Oregon State University, "An Analysis of Analysis". January, 2012.
- NIPS BigLearning: Workshop on Algorithms, Systems, and Tools for Learning at Scale, "Big Machine Learning Made Easy". December, 2011.
- Sandia Labs, "An Analysis of Analysis". June, 2011.
- Accenture Labs, "New Directions for Multi-label Classification". April, 2011.
- Kodak Digital Imaging Forum, "Submodularity and Sparse Coding". November, 2009.
- 2008 Digital Imaging Conference, "Machine Learning and Image Processing Techniques for Automatic Print Dating". November, 2008.
- Kodak Digital Imaging Forum, "Machine Learning Challenges in Kodak's GCG Business Unit". April, 2008.

- Kodak Digital Imaging Forum, "Four Short Lectures on Machine Learning". December, 2007.
- University of Rochester, "Structured Gradient Boosting". October, 2007.
- Kodak Intelligent Systems Research Center, "Machine Learning for Sequence Data Retrieval: A Comprehensive Approach". June, 2007.

REFEREEING ACTIVITIES

- Reviewer, *NIPS: Neural Information Processing*
- Program Committee, *AAAI Conference on Artificial Intelligence*
- Guest Editor, *Machine Learning Journal Special Issue on Structured Prediction*
- Reviewer, *International Conference on Artificial Intelligence and Statistics*
- Reviewer, *ACM Conference on Human Factors in Computing Systems*
- Reviewer, *Pattern Analysis and Machine Intelligence*
- Reviewer, *International Conference on Machine Learning*
- Reviewer, *International Conference on Image Processing*
- Organizing Committee Chair, *ICML 2007 Workshop on Constrained Optimization and Structured Output Spaces*
- Reviewer, *Machine Learning Journal*
- Reviewer, *International Symposium on Music Information Retrieval*

OTHER INTERESTS

- Freelance pianist, working independently and with several groups performing throughout the united states. Produced, directed, and performed in "Beatles at OSU", a benefit concert series raising over $30,000 for the Center Against Rape and Domestic Violence, 2006-2009.
- Visiting scientist for the Rochester Area Colleges Center for Excellence in Math and Science, engaging in classroom experiments to help provide excellence in science education for Rochester area middle school students, 2007-2011.
- Volunteer at *Bay Cliff Health Camp* for summers 1998-2004, helping disabled adults have a summer camp experience by assisting them with mobility and personal care.

REFERENCES

Dr. Francisco Martin
2851 NW 9th
Suite D, Conifer Plaza Building
Corvallis, OR 97330
541.602.6696
martin@bigml.com

Jeffery Tran
1508 N. Damen #B4N
Chicago, IL 60622
847.322.6295
jeffery.tran@gmail.com

Dr. Henry Kautz
Computer Studies Building, Rm. 709
University of Rochester Box 270226
Rochester, NY 14627
585.275.5671
kautz@cs.rochester.edu

Dr. Majid Rabbani
36 Wexford Glen
Pittsford, NY 14534
585.477.3722
majid.rabbani@kodak.com

Dr. Prasad Tadepalli
3069 Kelley Engineering Center
Oregon State University
Corvallis, OR 97330
541.737.5552
tadepall@eecs.oregonstate.edu

Dr. Carman Neustaedter
250 - 13450 102nd avenue
Simon Fraser University
Surrey, B.C., V3T 0A3 Canada
+1 778.782.9034
carman_neustaedter@sfu.ca